

# Using construction kits to explain possibilities (Construction kits generate possibilities) (DRAFT: Liable to change)

Aaron Sloman

School of Computer Science, University of Birmingham.

---

**Installed:** 19 Nov 2014 (Relocated 15 Dec 2014)

**Last updated:**

30 Aug 2015; 5 Sep 2015 (Modified title); 28 Jan 2016; 13 Mar 2016

20 Nov 2014; 23 Nov 2014; 24 Nov 2014; 2 Jan 2015; 23 Jul 2015 (minor changes)

This paper is

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html>

A PDF version (possibly slightly out of date) is available here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.pdf>

This file was previously located at a different address, which is now being used for a more complete discussion of requirements for construction kits able to account for biological evolution:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>

A partial index of discussion notes available from this web site is in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html>

---

## CONTENTS

- [Background to this document](#)
  - [A terminological preamble: uses of the verb "explain"](#)
  - [What is science? Beyond Popper and Lakatos](#)
  - [Why allowing non-falsifiable theories doesn't make science soft and mushy](#)
  - [Critics of the proposal, e.g. Stich](#)
  - [Familiar construction kits](#)
  - [Concrete and abstract construction kits](#)
  - [Explaining what's possible vs explaining what happens](#)
  - [Historical note on "How is X possible?"](#)
  - [Sufficient vs Necessary explanations of possibilities/impossibilities.](#)
  - [An application of the theory: testing competences](#)  
Added 19 Oct 2015
  - [Proving impossibility](#)
  - [Eventual publication of the paper, in 1976, then as book-chapter](#)
  - [REFERENCES](#)
  - [Hofstadter Review of 1978 Book](#)
  - [Stich Review of 1978 Book](#)
-

## Background to this document

This document began life as a response to criticisms of Chapter 2 of The Computer Revolution in Philosophy (1978) made by reviewers (including Douglas Hofstadter (Hofstadter 1980)) and Stephen Stich (Stich 1981).

Chapter 2 of the book -- now available online here -- claimed that explaining how something (or some class of things) is possible is a major function of science, and the rest of the book presented tentative examples illustrating how AI (including computational linguistics) advanced our ability to explain (and sometimes predict) possibilities, as theories in physics and chemistry had done previously. For example,

- Chapter 6 attempted to explain (only in outline) how a machine with a certain sort of mind could deal creatively with a complex, changing and locally unpredictable universe;
- Chapter 7 attempted to explain how valid reasoning could use non-Fregean forms of representation including maps and diagrams depicting physical machines;
- Chapter 8 attempted to explain how a young learner with various procedural mechanisms could begin to use the natural number series for a variety of tasks involving setting up one-to-one correspondences over time, or in spatial configurations;
- Chapter 9 explained in outline how visual perception could make use of a combination of bottom-up, top-down, and middle-out processing straddling a variety of structural domains (not necessarily based on the kind of 3-D to 2-D projection later proposed by Marr and others as essential to biological vision);
- Chapter 10 began to speculate about how the sorts of mechanisms discussed in the book and others being developed in computer science and AI might account for some aspects of human consciousness, including explaining how some aspects of an individual's mind may be inaccessible to that individual while others are not.
- The Epilogue, a late addition, explained (semi-seriously) in outline how super-intelligent machines might improve our planet by limiting the freedom of humans, whose morals, knowledge and intelligence, on the whole, leave much to be desired.
- The Postscript, another late addition, attempted to explain how a computer programming language, like human languages, could be a rich and powerful tool yet allow the specification of contradictions arising from useless non-terminating procedures analogous to the expression of Russell's Paradox (and others) in human languages, a point developed in more detail in Sloman (1971)

Many major past scientific advances were theories about what is possible and explanations of some possibilities in terms of more fundamental possibilities.

The chapter is included part of the new (slightly) revised, freely available, online electronic edition of the book assembled in 2015, here: <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2>  
Still being modified from time to time.

## A terminological preamble: uses of the verb "explain"

There are at least five different, though closely related, uses of the verb "explain", and I often slide between them, which may cause confusion.

1. **A person can explain something:** Einstein explained the perceived aberrations in the orbit of Mercury. The detective explained how the thief had entered the building.
2. **A theory can explain something:** "General relativity explained the curvature of the path of a

- photon".
3. **A type of physical object or mechanism can explain something:** "Ice on a road can explain car crashes"; "Chemical bonds explain the strength of polymers".
  4. **A particular object or configuration can explain something:** "Ice on the road explained the crash"; "A missing bolt explained the crash".
  5. **A mathematical fact can explain something:** "The fact that there were 9 people explained why they could not all be arranged in pairs". Compare: "The fact that there were 11 people explained why they could not be arranged in a rectangular array apart from an 11x1 or 1x11 array".

This is not intended to be a complete list.

In the discussion below of how construction kits explain possibilities I may sometimes slide between these usages. I hope the context will always make clear what exactly is being said.

**NOTE: Some related ideas (Added 13 Mar 2016)**

I have been aware for some time that there is an overlap between my ideas about the role of explanations of possibilities, as opposed to laws, in science and some of Stuart Kauffman's ideas, e.g. in his (1995). I have now found much greater and more explicit overlap with this 2016 paper: Longo, Montevil & Kauffman (2012)

## **What is science? Beyond Popper and Lakatos**

**NOTE:**

Part of this introductory section is shared between two documents:

This document (on explaining possibilities):

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html>

and a much larger document on construction kits, especially construction-kits produced by biological evolution: Sloman [DRAFT 2016,2017....]

Chapter 2 of Sloman (1978) was an attempt to extend the philosophy of science developed by Karl Popper (1934) (and revised/extended in his later publications) which distinguished between scientific statements or theories and non-scientific (e.g. metaphysical) statements. The former were required to be empirically falsifiable, since if no falsifiable empirical statement can be derived from a theory it was said to lack empirical content. Unfortunately this criterion has been blindly followed by many scientists who seem to be ignorant of the history of science. E.g. the ancient atomic theory of matter was not falsifiable, but was an early example of a deep scientific theory. Popper (unlike many scientists who promote 'falsifiability' as a criterion for scientific content) acknowledged that some metaphysical theories could be important precursors of scientific theories, but it is arguable that labelling them 'metaphysics' rather than 'science' is arbitrary, in view of their importance for science.

For more on the ancient atomic theory see:

<http://plato.stanford.edu/entries/democritus/#2>

<http://en.wikipedia.org/wiki/Democritus>

Popper's philosophy of science was extended by Imre Lakatos (1980), who proposed ways of evaluating competing scientific research programmes, based on the sorts of progress they made over time. This shifts the problem away from taking final decisions about which theory (or research programme) is best, allowing evidence to mount up in favour of one or another over time, though always allowing for the possibility that some new development will shift the balance of support. Emphasising evaluation over an extended period of time, Lakatos distinguished progressive and degenerating research programmes. Requirements were specified for deciding which of two progressive research programmes is better, though it is not always possible to decide while both are

being developed. The history of science shows that what appears to be a decisive victory (like Thomas Young's evidence of diffraction of light, which was taken to disprove Newton's particle theory of light) can later be overturned (e.g. when light was shown to have a dual wave-particle nature).

My motivation, in 1978, for extending the work of Popper and Lakatos was based on the observation that many important scientific discoveries are concerned with what is *possible*, e.g. types of plant, types of animal, types of reproduction, types of thinking, types of learning, types of verbal communication, types of thought, types of mathematical discovery, types of atom, types of molecule, types of chemical interaction, and types of biological information-processing (a category that subsumes several of the other types). Investigation of varieties of biological information processing and the mechanisms (especially construction-kits) that support them is main focus of the Turing-inspired Meta-Morphogenesis project, whose aims were first formulated during the year of Turing's centenary 2012:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

A separate paper [Sloman \[DRAFT 2016\]](#) discusses in detail the ways in which specifications of "construction kits" provided by nature (e.g. physics and chemistry), including the mathematical properties and generative powers of those kits and space-time, can play a central role in answering questions about how various evolutionary processes are possible, thereby explaining how the products of those processes, namely enormously varied forms of life and the many biological mechanisms they use, all directly or indirectly products of natural selection, are possible.

#### **NOTE ADDED 3 Mar 2016**

I now think that one of the deepest and most interesting examples of a scientist trying to explain how something is possible is Erwin Schrödinger's attempt to answer the question "What is life?" ([1944](#)).

---

The rest of this paper focuses on the special properties of explanations of possibilities and why the important ideas of Popper and Lakatos about the nature of science have to be extended to accommodate them.

## **Why allowing non-falsifiable theories doesn't make science soft and mushy**

**(Added 24 Dec 2014)**

Many scientists influenced by Popper think that the falsifiability requirement is essential to rule out empty theories: if there is no such constraint how can we rule out theories about invisible fairies at the end of every garden, and other fairy-tales, as non-scientific? Answers were proposed in my original paper and the 1978 chapter (see the [publication history below](#)).

The answer given in the paper and book chapter had several components, namely: A theory purporting to explain how various objects, states of affairs, or processes are possible, should:

1. Use modes of inference that are already well understood and reliable, including logical inference, arithmetical calculation, computer simulation, and diagrammatic reasoning of sorts used in geometry, engineering, etc., defended in [Chapter 7](#) of the book. Conclusions about what is possible derived from the theory should be derived using methods whose reliability as modes of inference is well established. (That would, for example, rule out attempts to derive the possibility of the existence of a god from events humans cannot yet explain.)

2. There should be a clear demarcation between what the theory does and does not explain. For example, it should not depend on what some group of individuals find convincing as a mode of reasoning.
3. The theory should be *general*, that is, it should explain many significantly different possibilities, preferably including some possibilities not known about before the theory was proposed. This criterion should be used with caution. Insofar as a theory generates some possibilities not yet established by actual instances, efforts should be made to find or create instances. The more clues the theory gives as to where to look for instances the better: but this does not require the theory to make falsifiable predictions. If repeated efforts to find actual instances fail, this does not disprove the theory, but it does reduce its credit. (Compare the proposals by Lakatos for evaluating competing scientific research programmes.) So a theory should not "explain" arbitrarily large collections of possibilities, e.g. the possibility of new isotopes with prime numbers of neutrons being discovered on dates whose year and day are prime numbers, e.g. 17th Jan 2039.
4. The theory should explain "*fine structure*". That is descriptions of what can occur or exist that are derivable from the theory should be as rich and detailed as possible. Thus a theory merely explaining the possibility of different chemical elements in terms of different possible constituents of their atoms will not be as good as one which also explains how it is possible for the elements listed on the periodic table to have exactly the similarities and differences of properties implied in the table. (Developments in computer modelling techniques and the increasing power of computers allow explanations of new possibilities to be derived with far greater diversity and precision than in the past. Compare Turing's use of computer models in connection with his ideas about morphogenesis in biology [Turing \(1952\)](#).)
5. The theory should be *non-circular*, i.e. the possibilities assumed in the theory should not be of essentially the same character as the possibilities the theory purports to explain. Many philosophical and psychological theories fail this test because they propose internal mechanisms described using concepts of ordinary language, presupposing competences of the sort being explained, whereas computer-based models of human competence can pass the test, since assuming the possibility of information processing machinery (e.g. something like digital computers, or rule-interpreters, or neural nets) is quite different from assuming the possibility of a mind! However, notice that a kind of circularity, namely recursion, is possible *within* such an explanation. (Behaviourist psychology is partly based on a failure to see this.) Compare the use of [the designer stance](#) with the "intentional stance" when proposing theories about minds.
6. The derivations from the theory should be *rigorous*, i.e. within the range of possibilities explained by the theory, the procedures by which those possibilities are deduced or derived should be explicitly specified so that they can be publicly assessed, and not left to the intuitions of individuals. If the theory is very complex, the only way to find out exactly what it does and does not imply (or explain) may be to express it in a computer program and observe the output in a range of test situations. (This takes the place of logical or mathematical deduction.) In fact such rigour is very rarely achieved in the human and social sciences, though the use of computer models has made a large difference. (Of course, use of a computer model that leads to correct predictions is not in itself a proof that the theory on which the model is based is correct.)
7. The theory should be *plausible*: that is, insofar as it makes any assertions or has any presuppositions about what is the case or what is possible, these should not contradict any known facts. However, sometimes the development of a new theory may lead to the refutation of previously widely held beliefs, so this criterion has to be used with great discretion.

8. The theory should be *economical*: i.e. it should not include assumptions or concepts which are not required to explain the possibilities it is used to explain. Of two theories T1 and T2 purporting to explain how X is possible, if T1 makes more assumptions than T2, then T2 can provisionally be judged the better theory. However, future evidence could switch the verdict.

Often economy in science is taken to mean the use of relatively few concepts or assumptions, from which others can be derived as necessary. This is not always a good thing to stress, since great economy in primitive concepts can go along with uneconomical derivations and great difficulty of doing anything with the theory, that is, it can go along with *heuristic poverty*. For instance, the logicist basis for mathematics proposed by Frege, Russell and Whitehead is very economical in terms of primitive concepts, axioms, and inference rules, yet it is very difficult for a practising mathematician to think about deep mathematical problems if he expresses everything in terms of that basis, using no other concepts. Replacing numerical expressions by equivalents in the basic logical notation produces unmanageably complex formulae, and excessively long and unintelligible proofs. The main points get buried in a mass of detail, and so cannot easily be extracted for use in other contexts. More usual methods have greater heuristic power. So economy is not always a virtue. This is also true of Artificial Intelligence models.

9. The theory should be rich in *heuristic power*: i.e. the concepts, assumptions, symbolisms, and transformation procedures of the theory should be such as to make the detection of gaps and errors, the design of problem-solving strategies, the recognition of relevant evidence, and so on, easily manageable. This is a very difficult concept to define precisely, but it is not a subjective concept. The heuristic power of a theory may be a consequence of its logical structure, as people working in artificial intelligence have been forced to notice.

See [chapter 7 of CRP](#) and [McCarthy and Hayes, 1969](#), for more on this: "Some philosophical problems from the standpoint of Artificial Intelligence"

10. The theory should be *extendable* (compare Lakatos 1970). That is, it should be possible to embed the theory in an improved enlarged theory explaining more possibilities or more of the fine-structure of previously explained possibilities. For instance a theory explaining how people understand language, which cannot be combined with a perceptual theory to explain how people can talk about what they see, or use their eyes to check what they are told, is inferior to a linguistic theory which can be so extended. Extendability is a major criterion for assessing artificial intelligence models of human abilities. However, it is a criterion which often can only be applied in retrospect, after further research attempting to extend the model or theory. This could be called a requirement for a mechanism to 'scale out', as opposed to the more familiar requirement for mechanisms to 'scale up', namely deal with larger or more detailed problems. Scaling out requires close interaction with other mechanisms included in a theory to explain other phenomena. For example, at present (2014) no AI or neural theory of vision scales out adequately, e.g. by showing how vision plays a role in perception of possibilities, perception of constraints on possibilities and in mathematical understanding, as illustrated in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html>

So a good explanation of a range of possibilities should be definite, general (but not too general), able to explain fine structure, non-circular, rigorous, plausible, economical, rich in heuristic power, and extendable (allow scaling out). It is not at all easy to produce explanations of possibilities that meet all these requirements. [I think it can be shown that many highly regarded models/theories in AI, psychology, neuroscience, linguistics, and philosophy fail to meet them.](#)

## Critics of the proposal, e.g. Stich

Many readers found the chapter hard to make sense of, e.g. Stephen Stich in his helpful review of the book (Stich 1981). I now think it would have been useful to illustrate the idea of explaining a set of possibilities by making use of the notion of a *construction kit*. A construction kit, e.g. Lego, Meccano and others, includes a set of "primitive" components, with properties and relationships between them, which enable them to be combined into larger wholes, while their properties and relationships constrain the types of larger object into which they can be assembled.

Features of a meccano set can be used to explain how a particular sort of toy vehicle or toy crane can exist, by showing how each can be assembled from the parts available, subject to the constraints of the kit. E.g. melting down metal parts and then re-shaping them is ruled out. Likewise features of a grammar and vocabulary can be used to explain how a particular sentence is possible by showing how the sentence uses words in a particular lexicon assembled in accordance with the rules of the grammar.

A specification of the components and assembly processes of each of Meccano and Lego sets can be used to explain how each can produce structures that the other cannot produce. A process of construction suffices to demonstrate that something is possible using a particular kit, e.g. a particular type of toy crane. Something deeper is required in order to explain why something is impossible for a particular kit, i.e. not included in the set of possible constructs supported by the kit.

For example a toy crane with a jib that is hinged at the bottom and can be raised and lowered is possible using meccano. Demonstrating that it is not possible using only lego-bricks can make use of a number of facts, such as that the processes of assembly of those bricks always produce only rigid structures, since there is no hinge mechanism and now means of creating a hinge mechanisms. A related argument is that the process of assembling a structure using lego bricks constrains each brick added to have edges with exactly three orientations (i.e. all edges can be divided into three classes where edges in each class are parallel, and edges in different classes are perpendicular to each other).

The idea of spaces of possibilities generated by different sorts of construction kits may be easier for most people to understand than the comparison with generative powers of grammars mentioned in the chapter. The idea of a construction kit is also more directly relevant to a host of types of scientific explanation, as well as theories in engineering.

This is a first draft attempt to spell out that idea, and will be expanded later.

## Familiar construction kits

I hope most readers will be familiar with the fact that different sorts of things can be built with different sorts of construction kits, e.g. meccano, lego-bricks, tinker-toy, FischerTechnik, plasticine, sand, mud, paper and scissors, paper plus folding operations (origami) etc. More complex construction kits can sometimes be formed by combining simpler kits, e.g. combining Lego and meccano to produce a new composite construction kit. Additional hybrid parts might be required to improve the integration of the two styles.

For many scientific and engineering purposes we are interested not only in what can be built, but what the things that are built can do, e.g. how they can change shape, interact with other things, be extended, come apart, etc. etc.

Each kit, simple or combined, allows a space (domain?) of possible structures (and possible processes involving those structures). The spaces have different contents because of different mathematical features of the generating elements (parts and modes of composition).

Each explains possibilities in a domain without predicting which possibilities will be realised (which generally depends on external factors).

However there is an element of prediction insofar as the theory of a domain specifies constraints on ways in which complex instances can be extended.

For example, if you have already built some sort of structure using a kit, then, if the kit has not been exhausted, there will be alternative possible ways of extending that structure by adding one or more new parts. Each such extension will then normally remove some of the old possibilities for change and produce new possibilities for change.

To that extent there is predictive power in the theory of what the kit makes possible, though the predictions are not about [what will happen](#) after some change, but predictions about [how sets of possibilities will be altered](#).

[I think Immanuel Kant had more than an inkling of this.]

## Concrete and abstract construction kits

Construction kits for children include physical parts that can be combined in various ways to produce new physical objects that are not only larger than the initial components but have new shapes and new behaviours. Those are *concrete* construction kits.

There are also *abstract* construction kits such as grammars, axiomatic systems, computer programming languages, and programming toolkits. For more on varieties of construction kit see the discussion of concrete, abstract and hybrid construction kits [Sloman \[DRAFT 2016\]](#).

## Explaining what's possible vs explaining what happens

**Added: 23 Nov 2014. Modified: 22 Dec 2014**

Suppose someone uses a meccano kit to construct a toy crane, with a jib that can be moved up and down by turning a handle, and a rotating platform on a fixed base, that allows the direction of the jib to be changed. What's the difference between explaining how that is possible and how it was done? First of all, if nobody actually builds such a crane then there is no actual crane-building to be explained: yet, insofar as the meccano kit makes cranes like that possible it makes sense to ask how it is possible. This has at least two types of answer. (More types of answer are discussed in [the document on construction-kits](#).)

A1: The first answer is concerned with identifying the parts and relationships between parts that are supported by the kit, and how a crane of the sort in question could be composed of such parts arranged in such relationships.

A2: The second answer would describe a sequence of steps by which such a collection of parts could be assembled from the basic components provided by the kit. There may be many different sequences leading to the same result: identifying any one of them explains how the construction is possible, as well as how the end result is possible.



Both answers are correct, though A2 obviously provides more information than A1. Neither explanation presupposes that the possibility in question has ever been realised. This is very important for many engineering projects where something new is proposed and critics believe that the object in question could not exist, or could not be brought into existence using available known materials and techniques. The designer could answer sceptical critics by giving either an answer of type A1, or type A2, depending on the reasons for the scepticism. So from this point of view explanations of possibilities have much broader applicability than explanations of things that actually exist, since what actually exists is only tiny subset of what could possibly exist.

The associated document on [construction kits](#) subdivides explanations of type A2 into a variety of different sub-cases. (Work in progress.)

## **Historical note on "How is X possible?"** **(Added 23 Nov 2014)**

"How is X possible?" was a type of question raised for various cases of X by Immanuel Kant (e.g. how is knowledge of synthetic necessary truths possible?). In the early 1970s I wrote a paper about this, expanding on my 1962 DPhil thesis defending Kant's views of mathematical knowledge. The new paper attempted to show that claims about possibility and explanations of possibility are deeply connected with the most fundamental aims of science, and often require the current scientific ontology to be extended.

As far as I knew no philosopher of science had addressed such claims and explanations. Moreover they are counter-examples to many philosophical accounts of how scientific theories are, or should be, evaluated. E.g. the claim that X (or something of type X) is possible can never be refuted by experiment or observation. However it can sometimes be confirmed by observation of X, or of type X. So stressing the scientific importance of questions and theories about what is possible and how those things are possible required challenging major philosophies of science emphasizing prediction and refutation, including the work of two philosophers whom I greatly admired and had learnt from, Karl Popper and Imre Lakatos.

Moreover, explaining how X is possible seemed to be particularly relevant to some of the newest sciences, including theoretical linguistics, computer science, and artificial intelligence.

## **Sufficient vs Necessary explanations of possibilities/impossibilities.**

"X makes Y possible" does not imply that if X does not exist then Y is impossible, only that one route to existence of Y is via X. Other things can also make Y possible, e.g., an alternative construction kit. So "makes possible" is a relation of sufficiency, not necessity.

An exception could be a case where X is the - the Fundamental Construction Kit discussed in - since all concrete constructions must start from it (in this universe?). If Y is abstract, there need not be something like the FCK from which it must be derived. The space of abstract construction kits may not have a fixed "root". However, the abstract construction kits that can be thought about by physically implemented thinkers may be constrained by a future replacement for the Church-Turing thesis, based on later versions of ideas presented here. Although the questions about explaining possibilities arise in the overlap between philosophy and science (Sloman, 1978, Ch.2), I am not aware of any philosophers who explicitly address the theses discussed here, though there are examples of potential overlap, e.g. Bennett (2011); Wilson (2015).

## **An application of the theory: testing competences**

**Added 19 Oct 2015**

A possible application of these ideas to testing competences is presented in a separate paper:

The relevance of explanations of possibilities to assessing competences

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/assessing-competences.html>

## **Proving impossibility**

**Added 24 Nov 2014**

After reading an early draft, Jack Birner reminded me that in many cases where it is impossible to prove *empirically* that X is impossible, the impossibility can be proved *mathematically*, for example Euclid's theorem that there cannot be a largest prime number

[https://en.wikipedia.org/wiki/Euclid's\\_theorem](https://en.wikipedia.org/wiki/Euclid's_theorem) and Arrow's theorem about the impossibility of a voting system that simultaneously satisfies certain superficially desirable requirements

[https://en.wikipedia.org/wiki/Arrow's\\_impossibility\\_theorem](https://en.wikipedia.org/wiki/Arrow's_impossibility_theorem)

There are many more examples of proofs of impossibility thousands of years old: Ancient mathematicians proved that it is impossible for the sum of the interior angles of a planar triangle to differ from half a rotation (180 degrees), as discussed in:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html>

**(Extended: 30 Aug 2015)**

It is worth noting that some proofs of impossibility depend on a set of premisses that can be modified or extended without any mathematical error. For example, there are "standard" proofs that it is impossible in Euclidean geometry to trisect an arbitrary angle using straight edge and compasses. That proof depends on limitations on uses of compasses and straight edge in Euclidean geometry. However, anyone who understands those constraints can easily understand a slight extension to its constructions by permitting translations and rotations of a straight edge with two marked points, known as the "Neusis" construction (known to Archimedes) and with that extension trisection of an arbitrary planar angle can be proved to be possible, as shown in:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html>

Turing proved that it is impossible for any Turing machine (TM) to be able to take in the specification of any arbitrary TM and derive a proof that it will, or will not, halt. If you understand what prime numbers are you may be able to construct a proof that it is impossible for any number of the form  $7^m - 5^m$  where m is a non-zero positive integer to be divisible by 5. (I have deliberately invented a theorem that is too specialised to be in any mathematical text book. You may prefer to try to prove a more general version that doesn't mention any particular numbers.)

Of course, any proof of an impossibility is a proof of some necessity, and vice versa, since the impossibility of P is the same as the necessity of Not-P, and the impossibility of Not-P is the same as the necessity of P. However, proof of possibility has a different character, though both proofs of possibility and proofs of impossibility start from prior knowledge about a "domain" that is under discussion (e.g. possible 2-D shapes in a plane surface). How to explain or model human knowledge about such domains remains problematic. The claim that logical reasoning abilities suffice was challenged in [Chapter 7](#) of the 1978 book, and more recently in a collection of papers discussing various examples of mathematical reasoning, including this discussion of "Toddler Theorems" <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>, and others on this web site: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html>.

## Eventual publication of the paper, in 1976, then as book-chapter

After rejection by at least one philosophy of science journal, the paper arguing that the aims of science included discovering and explaining possibilities, was published as "What are the aims of science?" in 1976 in *Radical Philosophy*, now online [here](#):  
<http://www.radicalphilosophy.com/issues/013>

A slightly revised version was published soon after as Chapter 2 of my (messy) 1978 book *The Computer Revolution in Philosophy: Philosophy, science and models of mind* (now freely available online [here](#)). The revised paper on explaining possibilities was Chapter 2, available [here](#).

To my surprise, several readers whom I thought would share my views told me that they had found that chapter hard to understand, and could not see its relevance to the rest of the book, although I had previously been encouraged by approving comment from a theoretical physicist colleague with a philosophical background, who later went on to receive a Nobel prize for physics. He, and the founding editor of the *Radical Philosophy* Journal may, for all I know, be the only two people who understood and liked that chapter, apart from some of my students.

In November 2014, I stumbled across a 1981 review of the 1978 book, by Stephen Stich, which also made critical comments about Chapter 2, while approving of much else in the book -- though highly critical (and rightly so) of much of the style of presentation. The text of his review is available [here](#) (added 19th Nov 2014, with his permission).

This new document attempts to provide a clearer introduction to the idea of a set of possibilities and the concept of an explanation of how something is possible, based on the idea of a construction kit (e.g. Lego, Meccano, plasticine, paper+scissors, a programming language, and many more) as a generator of a set of possibilities. A first draft was written in November 2014, but it is likely to be clarified and extended later. The main idea is that the physical world provides a very powerful (mostly chemical) construction kit that was "used" by natural selection to produce an enormous variety of organisms on this planet, some of which have produced new sorts of construction kits as toys or as major engineering resources.

We still have much to learn about the powers of that construction kit, the details of how those powers came to be used for life on earth, and what sorts of potential it has that have not yet been realised.

Further discussion of requirements for such a construction kit and its powers can be found in a separate document still under active development: [Sloman \[DRAFT 2016\]](#)

This can be read as a contribution to metaphysics. A closely related document on 'Actual Possibilities' (published in 1996) is freely available online [here](#).

---

## REFERENCES

- Karen Bennett (2011) Construction Zone: No Hard Hat Required. *Philosophical Studies*, 154:79-104, DOI 10.1007/s11098-011-9703-8
- Erwin Schrödinger, 1944, *What is life?*, CUP, Cambridge.  
Extracts from this book, with some added comments, can be found here:  
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html>

- <http://www.radicalphilosophy.com/article/what-are-the-aims-of-science>, Aaron Sloman, (1976,) 'What are the aims of science?' *Radical Philosophy*, 13, pp. 7--17,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2>  
Aaron Sloman (1978)  
'What are the Aims of Science?' Chapter 2 of *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind*  
Whole book: <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/>  
including [Chapter 2](#).
- **Hofstadter Review of 1978 Book**  
BULLETIN (New Series) OF THE AMERICAN MATHEMATICAL SOCIETY  
Volume 2, Number 2, March 1980 (Copyright 1980 American Mathematical Society)  
*The computer revolution in philosophy: Philosophy, science and models of mind*  
by Aaron Sloman, Harvester Studies in Cognitive Science Humanities Press, Atlantic Highlands, N. J., 1978, xvi + 304 pp., cloth, \$22.50.  
Reviewed by Douglas R. Hofstadter  
<http://www.ams.org/bull/1980-02-02/S0273-0979-1980-14750-3/S0273-0979-1980-14750-3.pdf>  
**Comment**  
(The review rightly criticises some of the unnecessarily aggressive tone and throw-away remarks, but also gives a thorough assessment of the main ideas of the book.  
However, like many researchers in AI (and probably most in philosophy, including Stephen Stich, [in the review referenced below](#)), Hofstadter regards the philosophy of science in the first part of this book, e.g. [Chapter 2](#), as relatively uninteresting, whereas I still think (in 2015) that understanding those issues is central to understanding how human minds work as they learn more about the world and themselves. Some of my recent work is still trying to get to grips with those issues in the context of a theory of varieties of learning and development in biological and artificial systems, e.g. in connection with the EU funded [CoSy robotic project \(2004-8\)](#), followed by the EU funded [CogX project \(2008-12\)](#), but most importantly by the (unfunded) Turing-inspired Meta-Morphogenesis project begun late 2011:  
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>
- Stuart A. Kauffman (1995), *At home in the universe: The search for laws of complexity*, Penguin Books, London,
- Giuseppe Longo, Mael Montevil and Stuart Kauffman, (2012) No Entailing Laws, but Enablement in the Evolution of the Biosphere, *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO '12*, ACM, New York, NY, USA, pp. 1379--1392,  
<http://arxiv.org/abs/1201.2069v1>
- <http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1971-03>  
Aaron Sloman, 1971, Tarski, Frege and the Liar Paradox, in *Philosophy*, 46, 176, April, 1971, pp.133--147,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>  
Aaron Sloman [REVISED DRAFT Jan 2016]  
Construction kits for biological evolution  
(Including evolution of minds and mathematical abilities.)  
(Discussion note.)

- **Stich Review of 1978 Book**  
<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/stich-review-crp.html>  
(Available here with the author's permission)  
Stephen P. Stich (1981)  
Review: *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind* (1978) by Aaron Sloman, in *The Philosophical Review*, Vol. 90, No. 2 (Apr., 1981), pp. 300-307
- Alastair Wilson, Metaphysical Causation, (Technical Note), July, 2015, University of Birmingham, Department of Philosophy, <http://alastairwilson.org/files/mc120715.pdf>  
Video lecture: <https://www.youtube.com/watch?v=j211wKrtlxs>
- A. M. Turing, (1952), The Chemical Basis Of Morphogenesis, *Phil. Trans. R. Soc. London*, B 237, 237, pp. 37--72,